

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

APPLICATION FOR LETTERS PATENT

Email Filtering Methods and Systems

Inventor(s):

Steven T. Maller

ATTORNEY'S DOCKET NO. MS1-353US

1 **TECHNICAL FIELD**

2 This invention relates to email filtering methods and systems.

3

4 **BACKGROUND OF THE INVENTION**

5 Email systems facilitate the exchange of electronic mail over a network,
6 such as a LAN (local area network), WAN (wide area network), or public network
7 (e.g., Internet). Most people are familiar with traditional computer email systems,
8 which are constructed on a client-server model. The email system involves a
9 server-based mail program residing on the server computer to manage the
10 exchange of email messages over one or more networks and a client-based mail
11 program residing on the client to implement a mailbox that receives and holds the
12 email messages for a user. The client-based mail program also implements a
13 graphical user interface that enables the user to open and read mail, or to create
14 new mail messages.

15 Other email systems have evolved that are more focused on the server side
16 of the equation than on the client side. In these types of systems, email servers
17 provide the bulk of the functionality that a client sees when the client enters the
18 email environment. One such system is shown generally at 20 in Fig. 1 and
19 includes an email server system 22 having a processor 24 that is configured to
20 receive email messages from a sender and distribute copies of the email messages
21 to one or more recipients. A recipient storage location 26 is provided and includes
22 a plurality of storage locations that are dedicated to individual recipients, with
23 three exemplary dedicated storage locations being shown at 28, 30 and 32. The
24 email messages that are received by the server system are copied directly into the
25 dedicated storage locations (sometimes referred to as “inboxes”) for each

1 recipient. If one message is received with nine specified recipients, then nine
2 copies of the message are made and placed into nine separate dedicated recipient
3 storage locations.

4 In systems such as these, many of the functions that are traditionally
5 implemented by a client machine are implemented by the server system 22. For
6 example, email messages, records, passwords, user preferences, address lists, and
7 the like are all saved on the server side in storage locations that are dedicated for
8 each recipient. Recipients or clients must then typically log in with the server and
9 run a browser program that lets them work within the email system to read their
10 messages. Logging in with the server is typically accomplished over a computer
11 network such as the Internet, and through the use of a suitable web browser. The
12 email environment is generated by the server through the use of HTML or web
13 pages that present the recipient or client with a screen that looks like an email box.
14 All email messages are delivered using the web page format. Microsoft's Hotmail
15 service is an exemplary system.

16 Email is a tremendously fast and efficient way to send electronic messages.
17 One of the problems that plagues the efficient use of email is the growing presence
18 of unwanted and unsolicited emails. These emails are typically referred to as
19 "spam." Spam can include unsolicited commercial emails (UCE) or non-
20 commercial emails. Spam is a menace that clogs email systems, slows down
21 performance, and severely impacts the manner in which email services are
22 provided by an email server to its clients.

23 From an operational standpoint, UCE or spam can consume vast amounts
24 of disk space and can monopolize many other machine resources. For example, in
25 the Fig. 1 example, when an email message is received by server system 22, a

copy for each intended recipient is made and placed into the dedicated storage location for each recipient. In the illustrated example, processor 24 receives a spam message and makes a copy of the spam message for each of the intended recipients, i.e. recipients 28, 30 and 32. If a large number of recipients are specified by the sender, then a large amount of server memory can be consumed by replicating the message and placing it in each dedicated storage location for each specified recipient. The typical message has the following fields that are shown in the figure: "TO", "FROM", "CC", "BC", and "SUBJECT".

From a customer service standpoint, system administrators are often at a loss to combat the delivery of spam to their individual clients. This can and often does result in large numbers of complaints and bounced email messages. Additionally, customers often do not wish to even receive certain types of morally or otherwise offensive emails. Yet, because the spammers (those who promulgate spam) predominate, innocent clients continue to be bombarded with unwanted email messages. From a legal standpoint, valuable time and resources are wasted in pursuing spammers because of the various havoc they wreak on network systems. Needless to say, spam continues to plague those who are in the business of providing email services to clients.

Accordingly, this invention arose out of concerns associated with providing improved systems and methods for reducing the impact that UCE and spam has on email recipients.

SUMMARY OF THE INVENTION

Various embodiments of the invention address two critical problems that current email service providers face. First, there is the problem of maintaining

1 high levels of customer service when email server systems are inundated with
2 spam. Second, there is the problem of reducing the system-wide impact that spam
3 has on the email delivery system. Current embodiments are directed to
4 determining whether an email message is an unwanted bulk email message
5 without necessarily considering the message that is conveyed by any portion of the
6 email message. Through analyses of patterns of delivery of these email messages,
7 profiles are built that allow an email server to ascertain whether there is a
8 likelihood that any one particular email message constitutes an unwanted email
9 message. If an email message is determined to likely constitute an unwanted
10 email message, then memory-saving measures are implemented. In preferred
11 embodiments, one copy of the email message is saved at a central, shared location
12 that can be accessed by each of the intended recipients. This avoids having to
13 replicate the email message across the system for each of the recipients.

14

15 **BRIEF DESCRIPTION OF THE DRAWINGS**

16 Fig. 1 is a high level diagram of a prior art email server system.

17 Fig. 2 is a flow diagram that describes a process in accordance with one
18 embodiment of the invention.

19 Fig. 3 is a diagram of a computer system that can be used to implement
20 various embodiments of the invention.

21 Fig. 4 is a high level diagram of an email server system in accordance with
22 one embodiment of the invention.

23 Fig. 5 is a diagram of the Fig. 4 email server system.

24 Fig. 6 is a flow diagram that describes processing in accordance with one
25 embodiment of the invention.

1 Fig. 7 is a diagram of a desirability index in accordance with one
2 embodiment of the invention.

3 Fig. 8 is a diagram of a user interface in accordance with one embodiment
4 of the invention.

5 Fig. 9 is a flow diagram that describes processing in accordance with one
6 embodiment of the invention.

7 Fig. 10 is a diagram of an email delivery system in accordance with a
8 preferred embodiment of the invention.

9

10 **DETAILED DESCRIPTION**

11 **Overview**

12 Unwanted email messages can come in many different packages such as
13 unsolicited commercial email (UCE) messages that seek to have recipients invest
14 or spend money, or email messages that are not necessarily commercial in nature,
15 but nonetheless are undesirable because of their message content, e.g. obscene or
16 morally inflammatory email messages. Various embodiments of the invention
17 recognize that unwanted emails messages often typically have delivery patterns
18 that distinguish them from other wanted emails. The delivery patterns are often
19 independent of the messages conveyed by any of the content of the email
20 messages, e.g. independent of the content of the sender's address field, the subject
21 field, and the message field. That is, determining that an email message is
22 unwanted does not necessarily depend on reading an email message and
23 interpreting its content. For example, some of the patterns that unwanted email
24 messages exhibit are that they typically specify a large number of recipient
25 addresses, and a determinable percentage of the specified recipient addresses may

1 be invalid. Another pattern concerns the size of the email message. That is, many
2 unwanted email messages are larger than a determinable size. Over time, profiles
3 are built based upon the these and other delivery patterns. These profiles are then
4 used by the server to assess an incoming email message and determine whether it
5 is likely to constitute an unwanted email message. Without reading the content of
6 an email message it is difficult to ascertain with complete certainty whether it is in
7 fact unwanted. Yet, the profiles that are built can assign a degree of likelihood that
8 any given email message does or does not constitute an unwanted email message.

9 If an email message likely constitutes an unwanted email message, then one
10 or more preventative measures can be taken to minimize the impact that the email
11 message has on its intended recipients and on the server system. For example, the
12 server might redirect an incoming email message if it has determined that the
13 email message is likely to be unwanted. Such redirection might involve, for
14 instance, placing the email message in a specially-defined “folder” or storage
15 location, which has been created to receive such messages. As another example
16 useful in various server-side embodiments, if a server determines that an incoming
17 email message likely constitutes an unwanted email message, the server makes a
18 single copy of it and places it in a shared storage location. This shared location is
19 represented to the user as a special folder or storage location similar to an inbox.
20 Individual recipients are notified and can read the email message at the shared
21 location. Thus, the server avoids having to make and store numerous copies of the
22 email message. If a recipient so desires, they can request a copy of the email
23 message at which time the server places a copy of it in a dedicated recipient
24 storage location for them.

1 Embodiments of the invention also provide a robust collection of heuristics
2 that go far beyond simply screening an email message based upon a sender's
3 address. The heuristics can consider parameters that are independent of the
4 message conveyed by any content portion of an email message, and/or parameters
5 that are dependent upon the message thus conveyed. Flexibility is provided by
6 enabling the parameters to be adjusted to accommodate different patterns. For
7 example, one set of heuristics might be used during prime computing time, while
8 another set is used during slower computing times.

9 Further, solutions are provided that enable the server and its clients to work
10 in concert to uniquely tailor the server's screening to fit the individual needs of the
11 clients. The concept of a desirability index establishes values that are assigned to
12 various degrees of desirability that an email message can have. Various
13 parameters having parameter values are associated with each index value. A user
14 can adjust either the parameter values or the index values to establish a degree of
15 desirability. Email messages are then evaluated against the defined index value.

16 Fig. 2 shows a flow diagram at 100 that describes processing in accordance
17 with embodiments of the invention. The processing that is described preferably
18 takes place on the server side and is executed by the server system. The server can
19 be a dedicated server that is specifically programmed to screen email messages. A
20 profile of unwanted email is first developed at step 102. The profile preferably
21 takes into account information that is not dependent upon the message conveyed
22 by any of the content of an email message. The server is configured at step 104 to
23 screen email messages based upon the profile. The server evaluates each of the
24 email messages that it receives at step 106. If an email message meets the profile,
25 then the server places a copy of the email message in a central location at step 108

1 and sends a notification to each of the recipients that an email message has been
2 received. In preferred embodiments, only one copy is saved by the server,
3 although other copies can be made if a recipient has specifically requested to
4 receive email messages from a particular sender. The one copy is placed in a
5 location that can be shared by each of the intended recipients for reading the
6 message. If an email message does not meet the profile, then the server delivers
7 the email message to a recipient storage location at step 110. In this way, email
8 messages that have a high likelihood of constituting unwanted email or spam are
9 not replicated across the server's storage system. Rather, storage is conserved by
10 saving only one copy of the email message.

11

12 Computer System

13 Preliminarily, Fig. 3 shows a general example of a desktop computer 130
14 that can be used in accordance with the invention. Computer 130 can be used to
15 implement server or client machines.

16 Computer 130 includes one or more processors or processing units 132, a
17 system memory 134, and a bus 136 that couples various system components
18 including the system memory 134 to processors 132. The bus 136 represents one
19 or more of any of several types of bus structures, including a memory bus or
20 memory controller, a peripheral bus, an accelerated graphics port, and a processor
21 or local bus using any of a variety of bus architectures. The system memory 134
22 includes read only memory (ROM) 138 and random access memory (RAM) 140.
23 A basic input/output system (BIOS) 142, containing the basic routines that help to
24 transfer information between elements within computer 130, such as during start-
25 up, is stored in ROM 138.

Computer 130 further includes a hard disk drive 144 for reading from and writing to a hard disk (not shown), a magnetic disk drive 146 for reading from and writing to a removable magnetic disk 148, and an optical disk drive 150 for reading from or writing to a removable optical disk 152 such as a CD ROM or other optical media. The hard disk drive 144, magnetic disk drive 146, and optical disk drive 150 are connected to the bus 136 by an SCSI interface 154 or some other appropriate interface. The drives and their associated computer-readable media provide nonvolatile storage of computer-readable instructions, data structures, program modules and other data for computer 130. Although the exemplary environment described herein employs a hard disk, a removable magnetic disk 148 and a removable optical disk 152, it should be appreciated by those skilled in the art that other types of computer-readable media which can store data that is accessible by a computer, such as magnetic cassettes, flash memory cards, digital video disks, random access memories (RAMs), read only memories (ROMs), and the like, may also be used in the exemplary operating environment.

A number of program modules may be stored on the hard disk 144, magnetic disk 148, optical disk 152, ROM 138, or RAM 140, including an operating system 158, one or more application programs 160, other program modules 162, and program data 164. A user may enter commands and information into computer 130 through input devices such as a keyboard 166 and a pointing device 168. Other input devices (not shown) may include a microphone, joystick, game pad, satellite dish, scanner, or the like. These and other input devices are connected to the processing unit 132 through an interface 170 that is coupled to the bus 136. A monitor 172 or other type of display device is also

1 connected to the bus 136 via an interface, such as a video adapter 174. In addition
2 to the monitor, personal computers typically include other peripheral output
3 devices (not shown) such as speakers and printers.

4 Computer 130 commonly operates in a networked environment using
5 logical connections to one or more remote computers, such as a remote computer
6 176. The remote computer 176 may be another personal computer, a server, a
7 router, a network PC, a peer device or other common network node, and typically
8 includes many or all of the elements described above relative to computer 130,
9 although only a memory storage device 178 has been illustrated in Fig. 2. The
10 logical connections depicted in Fig. 2 include a local area network (LAN) 180 and
11 a wide area network (WAN) 182. Such networking environments are
12 commonplace in offices, enterprise-wide computer networks, intranets, and the
13 Internet.

14 When used in a LAN networking environment, computer 130 is connected
15 to the local network 180 through a network interface or adapter 184. When used
16 in a WAN networking environment, computer 130 typically includes a modem 186
17 or other means for establishing communications over the wide area network 182,
18 such as the Internet. The modem 186, which may be internal or external, is
19 connected to the bus 136 via a serial port interface 156. In a networked
20 environment, program modules depicted relative to the personal computer 130, or
21 portions thereof, may be stored in the remote memory storage device. It will be
22 appreciated that the network connections shown are exemplary and other means of
23 establishing a communications link between the computers may be used.

24 Generally, the data processors of computer 130 are programmed by means
25 of instructions stored at different times in the various computer-readable storage

1 media of the computer. Programs and operating systems are typically distributed,
2 for example, on floppy disks or CD-ROMs. From there, they are installed or
3 loaded into the secondary memory of a computer. At execution, they are loaded at
4 least partially into the computer's primary electronic memory. The invention
5 described herein includes these and other various types of computer-readable
6 storage media when such media contain instructions or programs for implementing
7 the steps described below in conjunction with a microprocessor or other data
8 processor. The invention also includes the computer itself when programmed
9 according to the methods and techniques described below.

10 For purposes of illustration, programs and other executable program
11 components such as the operating system are illustrated herein as discrete blocks,
12 although it is recognized that such programs and components reside at various
13 times in different storage components of the computer, and are executed by the
14 data processor(s) of the computer.

15

16 **Email Server System Architecture**

17 Fig. 4 shows an exemplary email server system 34 in accordance with one
18 embodiment of the invention. System 34 includes one or more email servers 36
19 and a recipient storage location 38. Server 36 includes an email filter or filter
20 processor 40 having a heuristic library 42 and a storage location 44 that is
21 managed by email server 36. Although email server system 34 is shown as an
22 integral unit, it is to be understood that the various constituent parts thereof can be
23 separately implemented and associated with one another. Recipient storage
24 location 38 includes a plurality of dedicated storage locations for each of the users
25 or recipients, with exemplary locations being shown at 46, 48 and 50. Bulk email

1 messages are received at the email server location and are typically addressed to a
2 plurality of recipients. Server 36 filters various email messages by applying one
3 or more heuristics that are defined in heuristic library 42 to an incoming email
4 message. The heuristics enable the filter to determine whether an email message
5 likely constitutes an unwanted email message. One of the goals of some
6 embodiments of the present invention is to ascertain whether an email message is
7 unwanted without having to consider the message that is conveyed by the text of
8 the email message. Accordingly, some of the heuristics are directed to looking at
9 aspects of an email message that are not necessarily related to the content of the
10 message itself. Thus, some heuristics consider the number of recipients to which a
11 particular email message is addressed and the overall size of the email message. If
12 the email message likely constitutes an unwanted email message, it is redirected to
13 a first location such as storage location 44, rather than being placed in one or more
14 of the dedicated storage locations 46, 48 and 50 for the addressed recipients. This
15 constitutes but one way of ensuring that the server does not make as many copies
16 of the email message as there are specified recipient addresses. Email messages
17 that are found not likely to constitute unwanted email messages are delivered to
18 the individual respective dedicated storage locations 46, 48 and 50 and can be
19 provided to the recipients in due course.

20 In a preferred embodiment, only a single copy of the email message is
21 maintained at storage location 44. Other copies can be made if a particular
22 recipient has requested to receive email messages from a particular sender, as will
23 become apparent below. This avoids having to reproduce the email message for
24 each and every addressed recipient and store the email messages at multiple
25 locations throughout the server system. Since many unwanted email messages are

1 bulk in nature (i.e. addressed to many intended recipients) storage location 44 can
2 also be considered as a bulk email folder into which bulk email messages can be
3 first placed.

4 Fig. 5 shows the Fig. 4 system after an email message has been received
5 and determined to constitute an unwanted bulk email message or spam. The spam
6 message is stored at storage location 44. Server 36 then notifies each of the
7 intended recipients or specified addressees that an email has been received for
8 them. One way of sending notification to the recipients is to place a pointer at a
9 second location that is dedicated to each recipient, e.g. an email folder of each
10 recipient. Then, when the recipient logs in to check their email, the pointer will
11 allow them to access and read the stored email message if they so desire. This is
12 preferably done without making any copies of the email message, other than the
13 one that is stored at storage location 44. The recipient is also free to request that
14 the server make a copy of the email message and store the email message in a
15 dedicated recipient-specific storage location. In the illustrated example, the third
16 recipient has requested a copy of the spam message which is now stored in their
17 own dedicated storage location. In this example, storage location 44 is shared
18 among the intended recipients. However, the fact that it is shared is virtually
19 transparent to the recipients. By using the pointer that is placed at the second
20 location, each recipient can open the corresponding email message (in which case
21 they can view the shared copy), delete the email message (in which case the
22 pointer is deleted), or move the email message (in which case an actual copy of the
23 email message is itself generated and stored in a location designated by the
24 recipient).

1 Fig. 6 shows a decision diagram generally at 200 that describes decision-
2 making that can take place in accordance with an embodiment of the invention.
3 Step 202 determines whether an email message matches a pattern that is associated
4 with spam or unwanted email messages. One way of defining patterns is through
5 the use of heuristics that are discussed below in more detail. If the email message
6 likely constitutes an unwanted email message, then step 204 delivers one copy of
7 the email message to a central location such as location 44. Preferably, the central
8 location is one that can be shared by the intended recipients to read the email
9 thereby eliminating the need to reproduce the email system wide for all of the
10 intended recipients. Step 206 notifies the intended recipients that an email
11 message has been received. The recipients are then free to access the email
12 message at the shared location and request that they receive their own copy of the
13 email message. If step 202 determines that an email message is not likely to be
14 unwanted, then step 208 delivers the email message to a recipient location, such as
15 locations 46, 48 and 50 in Fig. 5.

16

17 **Heuristics**

18 One of the advantages of the present invention is the robust collection of
19 heuristics that can be used by email server 36 to screen for unwanted email.
20 Preferably, the heuristics are built upon the principle that unwanted email or spam
21 typically exhibits a pattern when it is delivered. Many times this pattern is
22 independent of the message that is conveyed by any of the content of a particular
23 message. For example, one need not necessarily read the content of an unsolicited
24 email describing adult web sites if it can be ascertained from other aspects of the
25 message that it is likely to constitute unwanted email. For example, unwanted

1 email is typically addressed to a large number of specified recipient addresses.
2 These addresses may not appear in the "To:" field, but rather will appear in the
3 "BC" (blind copy) field, as in Fig. 1. Many times, a number of these addresses
4 will be invalid as a result of being computer generated to cover a large number of
5 address permutations. Additionally, some unwanted email messages might be
6 very large in size. By recognizing these aspects of an email message, heuristics
7 are designed that consider factors unrelated to a message's content. Content-based
8 filtering can, however, still be used in connection with one or more of the
9 embodiments of the invention. Hence, determinations concerning whether an
10 email message is likely or not to constitute spam can be made, in some
11 embodiments, without accessing any content of the sender's address field, the
12 subject field, or the message field.

13 As a simple example only, nine heuristics are set forth in Table 1 below.
14 These heuristics are not intended to limit the invention in any way. Rather, they
15 are only given to illustrate certain approaches that can be taken.

16

17 **Table 1**

18 Heuristic	19 Heuristic Description
20 1	Is the email message addressed to more than <insert number> recipients?
21 2	Is the email message addressed to more than <insert number> percent of invalid addresses?
22 3	Is the email message larger than <insert number> bytes?
23 4	Is the email message indirectly addressed to more than <insert number> recipients?
24 5	Is the email message delivered after 11:30 P.M local time?
25 6	Heuristic 5 and any of Heuristics 1, 2, 3, or 4.
7	Heuristics 1 and 2.

1	8	Heuristics 1 or 2.
2	9	Heuristics 1 and 3.

3
 4 Heuristic 1 determines whether a particular email message is addressed to a
 5 definable number of recipients. This recognizes a pattern that certain email
 6 messages that are unwanted are often addressed to a large number of recipients.
 7 The heuristic includes a parameter that is indicated by the “<insert number>” field.
 8 This permits the heuristic to be adjusted to accommodate different delivery
 9 patterns. Heuristic 2 determines whether a certain percentage of the specified
 10 addresses are invalid. This recognizes a pattern that often times an unwanted
 11 email message will be addressed to a large number of invalid accounts. The
 12 heuristic also includes a parameter that is indicated by the “<insert number>” field
 13 that permits the heuristic to be adjusted. Heuristic 3 determines whether an email
 14 message is larger than a definable size and includes a parameter that is indicated
 15 by the “<insert number>” field. This permits the heuristic to be adjusted to
 16 accommodate different delivery patterns. This heuristic recognizes a pattern that
 17 certain unwanted emails may have size characteristics that are distinguishable over
 18 ordinary email messages. Heuristic 4 determines whether an email message is
 19 indirectly addressed to a certain number of recipients and includes a parameter that
 20 is indicated by the “<insert number>” field. This permits the heuristic to be
 21 adjusted to accommodate different delivery patterns. This heuristic recognizes a
 22 pattern that unwanted email messages may have a large number of recipients
 23 “blind copied”. An email message can be considered as indirectly addressed to a
 24 recipient if the recipient’s address is not specified in the “TO” field, i.e. it appears
 25 either in the “CC” or “BC” fields. Heuristic 5 recognizes a pattern that certain

unwanted emails might typically be delivered after certain times in the evening.
The remaining heuristics constitute combinations of the heuristics mentioned
above.

The heuristics constitute but one way of defining different profiles that can be used to assess whether a particular email message is likely to be unwanted. By virtue of the fact that different combinations of heuristics can be used, and certain parameters values within certain heuristics can be varied, a robust set of flexible, adaptable profiles can be built and maintained. In addition, the profiles can be quickly adapted, system-wide, to address subtle changes in the delivery patterns of the spam.

Desirability Index

In one embodiment, the concept of a desirability index is used to assess email messages. Fig. 7 shows one such exemplary index at 300. The idea behind the desirability index is that index values, here 1-7, are assigned to various degrees of desirability that an email message can have. The degrees of desirability range from a low desirability value of 1 to a high desirability value of 7. The index values are associated with a plurality of parameters having parameter values. For exemplary purposes only, Table 2 sets forth the index values that are cross-referenced against some example parameters.

Table 2

Index Values	Number of specified recipient addresses	Percentage of invalid specified recipient addresses	Larger than X bytes	Delivery time
1	>1000	>20%	>X	Between 11:30 P.M. and 3:30 A.M.
2	0 < y <= 200	>10%	>X	Between 10:00 P.M. and 12:00 P.M.

1	3	$0 < y \leq 150$	5-15%	>X	Daytime
	4	$0 < y \leq 100$	5-10%	<X	Daytime
	5	≤ 30	0-10%	<X	Daytime
2	6	≤ 20	0-5%	<X	Daytime
	7	≤ 20	0-3%	<X	Daytime

3
4 The parameters in this example include: the number of specified recipient
5 addresses, the percentage of invalid specified recipient addresses, a size parameter,
6 and a delivery time parameter. The parameters each have values that correspond
7 to the various index values. Some of the parameters do not depend on any
8 message conveyed by any content of an email message. The parameter values are
9 preferably adjustable so that different patterns of delivery can be examined.

10 Fig. 8 shows a user interface 302 that can be used in connection with
11 desirability index 300. The user interface 302 is established so that a user, client,
12 or recipient can adjust either or both of the individual parameter values or the
13 index values. If the user adjusts a parameter value, then the index value associated
14 with a certain degree of desirability is made either more or less restrictive. If the
15 user adjusts the index value, then the user changes the degree of desirability. The
16 email server then uses the selected index value to assess and evaluate incoming
17 email messages for the user.

18 For example, when an email message is received at the server location, a
19 score can be calculated based upon one or more of the parameters. Any number or
20 combination of parameters can be used. In addition, parameters other than those
21 specifically shown can be used. The score is then compared with an index value
22 that is selected by a user or recipient. In this manner, the user-selected index value
23 represents a threshold value. The index or threshold value defines a likelihood
24 that a particular email message will constitute an unwanted email message. If an
25 email message's score exceeds the threshold value (here, in the negative

1 direction), then the email message likely constitutes one that a user or recipient
2 does not want. If this is the case, the server can then place a copy of the email
3 message at storage location 44 (Fig. 5) and send notifications to the intended
4 recipients.

5

6 Address Screening of Bulk Mail

7 Some bulk email messages, by their very nature, are desirable to some
8 users. By simply screening the bulk email messages as described above, it is
9 possible that some users might not see the bulk email messages. This is especially
10 so if the email messages at the single, shared location are only maintained for a
11 short determinable period of time. To address this situation, embodiments of the
12 invention provide additional address screening for bulk email messages. Address
13 screening permits a user or recipient to receive specified bulk email messages
14 without having to read the email messages from the single, shared location. The
15 user or recipient can specify a list of approved senders. This provides an
16 advantageous way for a user to receive bulk email messages that they do not want
17 to be filtered. For example, Microsoft may send bulk email messages that
18 describe various software upgrades to its users. A user who wishes to receive
19 these bulk email messages can simply add “microsoft.com” as a domain name
20 from which email messages will always be accepted. In addition, address
21 screening also permits users to screen bulk email messages based upon whether
22 they are directly addressed in the email message, i.e. their address appears in the
23 “TO” field and not in the “CC” or “BC” fields. This permits a user to screen
24 potential bulk email messages by looking for email messages in which they are
25 blind copied. Combining these two address screening techniques sets up a

powerful screening mechanism that allows only those bulk email messages that meet the specified criteria to be placed into a user-dedicated storage location. It should be apparent that as to other users who have not specifically placed a sender in their list of approved senders, the email message will be placed in the single, shared location and subsequent notifications will be sent out.

Fig. 9 shows a flow diagram at 400 that describes processing in accordance with one address screening embodiment. The server receives an email message at step 402 and determines it to be a bulk email message. Exemplary ways of determining whether an email message is a bulk message are described above and include determining the number of users or recipients to which the email message is addressed. The server then determines whether the bulk email message is addressed directly to a recipient at step 404. If a recipient's address appears in the "CC" or the "BC" field, it is not directly addressed to a recipient and is sent to the bulk email folder at step 408 and notifications are sent at step 410. If the email message is directly addressed to a recipient, step 406 determines whether the sender is on the recipient's list of approved senders. If the sender is not on the approved list, then the email message is sent to the bulk email folder at step 408 and notifications are sent at step 410. If the sender is on the recipient's list of approved senders, then the email message is delivered or placed in a dedicated recipient location at step 412, such as the user storage locations mentioned above.

Content Screening

Other embodiments of the invention recognize the fact that bulk email messages or spam often contains text that is identical or very similar to other bulk email messages or spam. This is the case, for example, when the same email

1 messages are sent at different times, or the same or similar email messages are
2 sitting in a queue awaiting delivery. For example, spammers may target some
3 recipients on one day, and then target other recipients on another day with the
4 same email message. If, for some reason, the spam escapes filtering when it is
5 first sent, then it may be possible to pick it up on the second day if it is similar in
6 content. Thus, a profile can be developed of unwanted email messages based
7 upon whether the email messages are similar in content with other email
8 messages. Email messages can be similar in content if they are identical or if they
9 contain a definable amount of textual similarities. The email server then looks for
10 email messages that meet the established profile. If the profile is met, then
11 processing can take place as described above.

12 Application of the profile can take place by conducting simple text searches
13 of the email messages that are received to ascertain whether any of their text
14 matches text of any other email messages such as those that have been found to be
15 spam. Alternately, different text-matching algorithms can be employed such as
16 fuzzy text-matching algorithms that impart a degree of intelligence to the email
17 server.

18

19 Hotmail Architecture

20 Preferred embodiments are implemented in conjunction with Microsoft's
21 Hotmail service. Fig. 10 is a diagram that shows general architectural features of
22 the Hotmail system generally at 500. A Hotmail "cloud" 502 encapsulates all of
23 the Hotmail functionality so that everything a user or client sees on their screen is
24 generated inside the Hotmail cloud. Hotmail cloud 502 includes an array of web
25 servers 504. When a user logs in, they communicate with one of the servers of the

array. The web servers are configured to serve web pages and do not contain any user data. The web servers can either pull HTML files off a storage disk or run a program to generate an appropriate HTML file. The file is then provided to a user browser executing on a user machine that requested the HTML file and is assembled by the browser at the user machine.

Hotmail cloud 502 includes one or more user database servers 506. All user or recipient data resides on the user database servers. This includes, for each account, all email messages, contact lists, personal preferences, passwords, and all other items typically associated with an email account. In practice, the user database servers are implemented by SUN Ultra Enterprise 4500-class servers. Each server stores from between 750,000 to 2,000,000 user accounts.

A database server 508 is provided and is an internal database server. Server 508 includes a list of all Hotmail users in memory, as well as the location of their user data on database servers 506. When a user contacts the Hotmail cloud 502, a web server of the web server array 504 contacts database server 508 to ascertain the location of the user's data on one of the user databases 506. The database server 508 returns the location to the web server which then can either assist a user in reading their email messages or assist a user in sending email messages.

When an email message is read by a user, the list of email messages are pulled by a web server of web server array 504. An appropriate web page is generated to appear as an email inbox. Links are embedded in the web page for the particular email messages. The web page is then sent to the user and assembled by the user's browser. The links retrieve the particular email messages for a user.

When email is sent, a user clicks on an appropriate composition page which brings up a web page that looks like an email page. The user types a message and clicks send. The email message is packaged as an *http* web request that is received by a server of the server array 504. The web server then contacts database server 508 to ascertain the location of the intended recipient. If the recipient exists, then their location is returned to the web server which then deposits the email message in the appropriate account. This process is utilized for the users that are within the Hotmail cloud 502. That is, these users are subscribers to the Hotmail email service. Email messages can, however, be received into the Hotmail cloud from outside of the cloud, e.g. from the users that are depicted outside of the cloud in Fig. 10.

Hotmail cloud 502 also includes an array of SMTP mail servers 510 that perform essentially the same as server array 504. That is, when an email message is received from outside of the cloud, a server of array 510 contacts database server 508 to ascertain a recipient location among user databases 506, and then deposits the email message at one or more of the appropriate locations. The SMTP servers are completely open to the network through which the email messages come. The servers are unable to authenticate whether the sender of an email message is an actual sender or not. Additionally, the SMTP servers 510 are unable to authenticate the server that connects with it to deliver the email message. Because of this, spam can easily enter the Hotmail cloud 502.

When spam arrives, the server array (either of arrays 504 and 510) makes as many copies of the spam as there are valid recipients, and places each copy in a dedicated user storage location provided by user database 506. With millions of

1 Hotmail users, it is easy to see how spam can be promulgated throughout the
2 Hotmail system and consume valuable and expensive memory.

3 In preferred embodiments, server arrays 504 and 510 are configured to
4 screen email message based upon a set of heuristics that determine whether an
5 email message is likely to constitute an unwanted email message. Alternately, a
6 dedicated server that is linked with each of these server arrays can perform the
7 email screening functions. Exemplary heuristics are discussed above. If an email
8 message is found to likely constitute an unwanted email message, a single copy is
9 placed in a storage location that is not a dedicated user storage location. The
10 server then notifies all of the valid specified recipients that an email message for
11 them has been received. This can take place as discussed above.

12

13 Conclusion

14 Various embodiments of the invention address two critical problems that
15 current email service providers face. First, there is the problem of maintaining
16 high levels of customer service when email server systems are inundated with
17 spam. Second, there is the problem of reducing the system-wide impact that spam
18 has on the email delivery system. Current embodiments are directed to
19 determining whether an email message is an unwanted bulk email message
20 without necessarily considering the message that is conveyed by any portion of the
21 email message. Through analyses of patterns of delivery of these email messages,
22 profiles are built that allow an email server to ascertain whether there is a
23 likelihood that any one particular email message constitutes an unwanted email
24 message. If an email message is determined to likely constitute an unwanted
25 email message, then memory-saving measures are implemented. In preferred

1 embodiments, one copy of the email message is saved at a central, shared location
2 that can be accessed by each of the intended recipients. This avoids having to
3 replicate the email message across the system for each of the recipients.

4 Although the invention has been described in language specific to structural
5 features and/or methodological steps, it is to be understood that the invention
6 defined in the appended claims is not necessarily limited to the specific features or
7 steps described. Rather, the specific features and steps are disclosed as preferred
8 forms of implementing the claimed invention.

9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25